REMARKS

The above amendments to the above-captioned application along with the following remarks are being submitted as a full and complete response to the Official Action dated December 6, 2007 and the telephone interview with the Examiners on April 1, 2008.

Applicants thank the Examiner and her supervisor for taking the time to conduct the telephone interview. The Examiner indicted that the proposed claim amendments appear to overcome the prior art of record; however, the Examiners deferred the consideration of the concurrently submitted IDS reference until receiving the formal response.

In view of the above amendments and the following remarks, the Examiner is respectfully requested to give due reconsideration to this application, to indicate the allowability of the claims, and to pass this case to issue.

Status of the Claims

Claims 1-22 are under consideration in this application. Claims 1, 8 and 15 are being amended, as set forth in the above marked-up presentation of the claim amendments, in order to more particularly define and distinctly claim Applicant's invention. A new claim 22 is being added.

The claims are being amended to correct formal errors and/or to better recite or describe the features of the present invention as claimed. All the amendments to the claims are supported by the specification. Applicant hereby submits that no new matter is being introduced into the application through the submission of this response.

Prior Art Rejections

Claims 1-2, 5-9, 12-16 and 19-21 were rejected under 35 U.S.C. §103(a) as being unpatentable over US Patent No. 5,555,371 to Duyanovich et al. (hereinafter "Duyanovich") over US Pat. No. 5,051,887 to Berger et al. (hereinafter "Berger"), US Patent No. 6,578,120 to Crockett et al. (hereinafter "Crockett '120") and US Pub No. 2001/0010070 of Crockett et al. (hereinafter "Crockett '070"), and claims 3-4, 10-11 and 17-18 were rejected under 35 U.S.C. §103(a) as being unpatentable over Duyanovich in view of Berger, Crockett '120, Crockett '070 and US Patent No. 6,938,045 to Li et al. (hereinafter "Li"). These rejections have been carefully considered, but are most respectfully traversed.

In a system (for example, the embodiment depicted in Fig. 1) including a plurality of primary storage subsystems 102 #1, 102 #2, a plurality of secondary storage subsystems 102 #3, 102 #4, 102 #5 that are connected to each other via a network SAN, and a host computer

101-2 including a remote copy manager (RCM) 105-2 ([0010]) and being connected with the secondary storage subsystems, a method of the present invention, as now recited in claim 1, remotely copies data from each of a plurality of primary volumes 110 directly via a remote copy link 104 to a corresponding secondary volume of a plurality of secondary volumes 117, wherein a number ("N"=2) of the plurality of primary storage subsystems is different from a number ("M"=3) of the plurality of secondary storage subsystems ("N:M configuration" [0004]), the primary volumes are constituted by the primary storage subsystems, and wherein the secondary volumes are constituted by the secondary storage subsystems. The method comprises (1) a normal synchronizing procedure (Figs. 4-5) and (2) an after-failure synchronizing procedure (Fig. 6-7).

(1) The normal synchronizing procedure includes: receiving, at each of the secondary storage subsystems, remote copy requests each of which is associated with a timestamp and directly sent from each one of the plurality of primary storage subsystems via the remote copy link 104 (Fig. 1; [0025]); receiving via the remote copy link 104, at each of the secondary storage subsystems, synchronizing requests each of which is associated with a timestamp and a primary storage ID of a primary storage subsystem, which sends a respective synchronizing request, from said each one of the primary storage subsystems respectively; determining, at each of the secondary storage subsystems, a first time as a first time parameter based on timestamps included in the synchronizing requests; and determining, at each of the secondary storage subsystems, which remote copy requests to process based on the first time parameter, primary storage IDs and timestamps associated with the remote copy requests, thereby maintaining data I/O consistency among said storage subsystems.

(2) The after-failure synchronizing procedure includes: suspending said remote copy requests from being directly sent from the primary storage subsystems via the remote copy link 104 to the secondary storage subsystems after a failure occurs, thereby starting a suspension period; collecting via a network link SAN (Fig. 1; [0042]) and comparing during the suspension period by the RCM in the host computer time parameters stored in the secondary storage systems to determine a synchronized time, said network link SAN being different from the remote copy link 104 and not overlapping with the remote copy link 104; receiving via said network link SAN during the suspension period from the RCM in said host computer at each of the secondary storage subsystems synchronizing requests each of which includes said synchronized time; updating during the suspension period second time parameters at each of the secondary storage subsystems up to said synchronized time; and determining during the suspension period, at each of the secondary storage subsystems,

which remote copy requests to process based on said updated second time parameter and processing a determined remote copy request by said each secondary storage subsystem therein, thereby maintaining data I/O consistency among said storage subsystems.

The invention recited in claim 8 is directed to software for remotely copying data from a plurality of primary volumes to a plurality of secondary volumes according to the method recited in claim 1.

The invention recited in claim 15 is directed to a system for remotely copying data from each of a plurality of primary volumes to a plurality of secondary volumes according to the method recited in claim 1.

The present invention improves over the Nano-copy technology ([0005]) to insure the consistency across multiple storage sub-systems with a N:M configuration (a N number of primary storage subsystems and a M ($\neq$ N) number of second storage subsystems). Data is asynchronously copied directly via RC Link 104 therebetween during a normal operation to do without management from a host ([0006]), i.e., not via any host/ data mover (XRC) as in Fig. 1 of Crockett '120. The present invention uses the RCM and via a network link SAN ***only*** during a suspension period (after-failure procedure).

During the normal operation, the primary storage subsystems of the present invention transmit update data in timestamps to the secondary system directly via RC Link 104 without any host intervention. This approach achieves data integrity in a significantly simpler and less costly hardware based solution without impact on server or application performance as much as applying the Nano-copy technology to a system of a N:M configuration ([0003]-[0005]).

During the after-failure operation (Figs. 6-7; defined in (2) part of claim 1), the RCM 101-2 at the secondary side of the preset invention collects and compares the time parameter of each secondary subsystem, while suspending said remote copy requests from being directly sent from the primary storage subsystems via the remote copy link 104 to the secondary storage subsystems ("a suspension period" of the RC link 104). In the claimed invention, the after-failure operation **overlaps** with suspension period of the RC link 104.

In contrast, the secondary SCU 31 in Crockett '120 executes the corresponding after-failure operation steps "**after**" the volume re-synchronization (during which communication is made between the primary SCU 27 and the secondary SCU 31 via the data mover 9 ; Fig. 7; col. 10, lines 19+), rather than "during the suspension period of the RC link." In particular, Crockett '120 merely suspends the primary volume 29 ("*the host CPU 2 will suspend the primary volume 29 ..... The SCUs in turn are still subject to application write updates, and they in turn will operate in a type of bit map mode. However, since the primary volume 29*

has been suspended, no record set copies of the write updates are made by the primary SCU 27." Col. 10, lines 49-54; p. 8, lines 11-12 of the outstanding Office Action), but not any link between the primary site 4 and the remote secondary site 6 as the present invention.

All the cited references show that the primary storage subsystem and the second storage subsystem are linked via a host/data mover or alike. If a host/data mover operates between the primary storage subsystems and the secondary storage subsystems, it is unnecessary to keep consistency of the stored data in N:M configurations with a direct link 104.

In the 1:1 configuration of Crockett '120 (Fig. 1), a primary site 4 communicates with a remote secondary site 6 over a CPU2 managed path including a channel extenders 19, 23 terminating a long haul telephone path 21, and a second ESCON channel 25 (Col. 7, lines 5-10). In the 1:1 configuration of Crockett '070 (Fig. 1), the data mover 114 serves to copy data from the primary storage 104 to the backup storage 106, thus mirroring the contents of the primary storage 104 at the backup storage 106 ([0031]). Berger's host CPU 110 is attached via its channels 120 to a controller 125 including dual DASD controllers 160 and 160' and attached to DASDs 70 and 75 (Fig. 2). Duyanovich's 1:1 configuration (Fig. 3) provides primary controllers 14 that reported a write operation to a data mover 104 (including time stamp and sequence number), and updates the data in primary data-storage system 13(Col. 13 lines 50-53), and also transfers data to the secondary system via another data mover 108 and secondary controllers 19. The cited references do not concern any N:M configuration, and would not use any the RC link between the primary and secondary systems as the present invention.

Even if, arguendo, the RC link 3 between the primary and secondary systems in Fig. 5 of the nanocopy reference were combined with the cited references, since the nanocopy reference fails to specify when the RC link 104 operates or suspended, one skilled in the art could not provide the method as claimed by the Applicants based on the above prior teachings except by using Applicants' invention as a blueprint. Applicants will point out that a rejection based on hindsight knowledge of the invention at issue is improper. Given the cited references together with Nano-copy, it is intuitive for one skilled in the art to use the data mover to execute the after-failure operation "after" Crockett's volume re-synchronization, since it is easier do carry out the after-failure operation this way given that the two remote sites already communicate via the data mover prior to the failure.

In addition, claim 22 recites that time parameters stored in the secondary storage systems are allocated at a fibre channel switch 901 in the SAN 103-1, rather than at the host computer ([0086]) as in the prior art.

Applicant contends that the cited references and their combinations all fail to teach or disclose each and every feature of the present invention as recited in independent claims 1, 8 and 15. As such, the present invention as now claimed is distinguishable and thereby allowable over the rejections raised in the Office Action. The withdrawal of the outstanding prior art rejections is in order, and is respectfully solicited.

Conclusion

In view of all the above, Applicant respectfully submits that certain clear and distinct differences as discussed exist between the present invention as now claimed and the prior art references upon which the rejections in the Office Action rely. These differences are more than sufficient that the present invention as now claimed would not have been anticipated nor rendered obvious given the prior art. Rather, the present invention as a whole is distinguishable, and thereby allowable over the prior art.

Favorable reconsideration of this application as amended is respectfully solicited. Should there be any outstanding issues requiring discussion that would further the prosecution and allowance of the above-captioned application, the Examiner is invited to contact the Applicant's undersigned representative at the address and telephone number indicated below.

Respectfully submitted,

Stanley P. Fisher
Registration Number 24,344

Juan Carlos A. Marquez
Registration Number 34,072

REED SMITH LLP
3110 Fairview Park Drive, Suite 1400
Falls Church, Virginia 22042
(703) 641-4200

April 4, 2008

SPF/JT